

# Intel® Array Building Blocks

**Dr.-Ing. Michael Klemm**  
Sen. Application Engineer  
Software and Services Group

# Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>



# Optimization Notice

## Optimization Notice

Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the “Intel® Compiler User and Reference Guides” under “Compiler Options.” Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3) instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors.

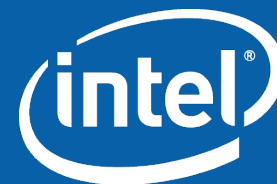
While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20101101



# Agenda

- **Overview and Goals**
- **How to add it to your project...**
- **Programming Constructs and Data Types**



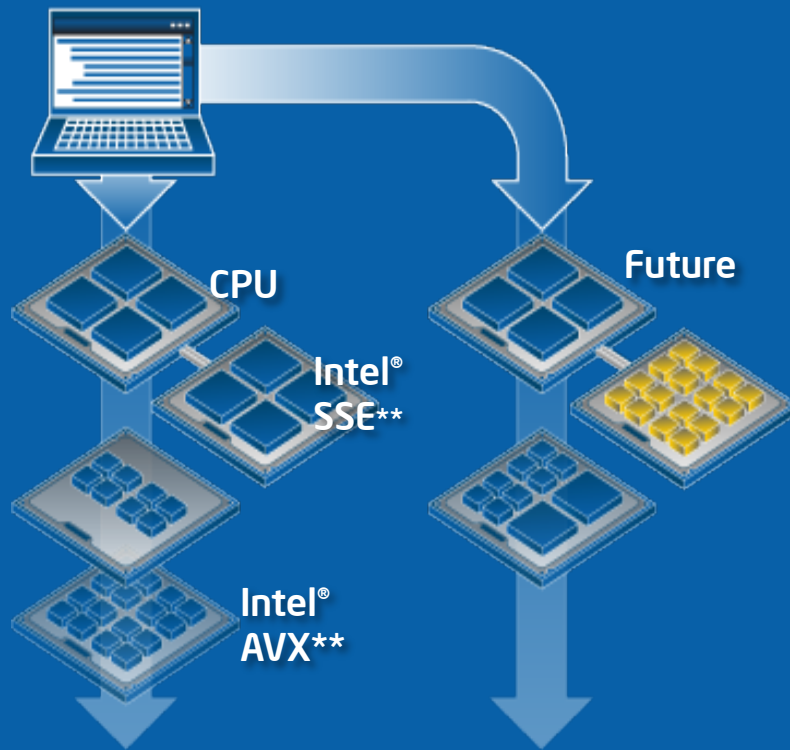
# Introduction to Intel® Array Building Blocks

Overview and Goals

# Introduction: Objectives

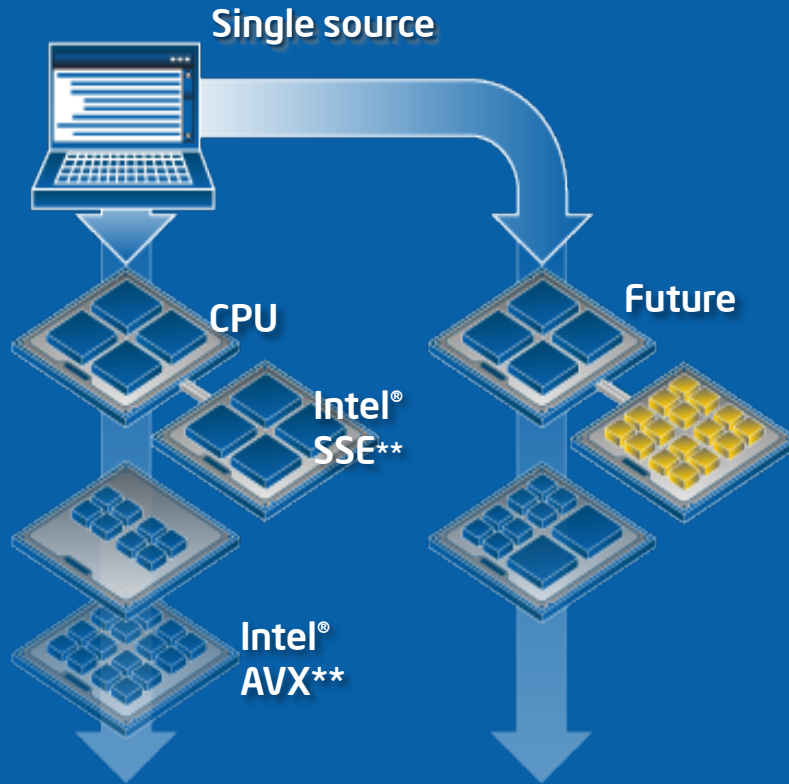
- **Understand the motivation for Intel® Array Building Blocks**
  - Also known as Intel® ArBB
- **Understand the Intel® ArBB C++ API-as-a-language**
- **Understand the basic syntax of the Intel® ArBB API**
- **Review the available operators**
- **Be able to write a first “Hello World” application w/ ArBB**
- **Work through a few example applications**

# What's Wrong with Parallel Programming?



- **Parallel programming is hard**
  - Deadlocks
  - Data races
  - Synchronization
  - Load imbalance
- **Errors inhibit productivity**
- **No uniform programming model for**
  - Intel SSE, Intel AVX
  - Multi-threading
  - IA manycore
- **Parallel programmers lose single code base for their applications**

# Intel® Array Building Blocks



## Productivity

- Integrates with existing tools
- Applicable to many problem domains
- Safe by default → maintainable

## Performance

- Efficient and scalable
- Harnesses both vectors and threads
- Eliminates modularity overhead of C++

## Portability

- High-level abstraction
- Hardware independent
- Forward scaling



# Productivity

- **Integrates**  
with existing IDEs, tools, and compilers: no new compiler needed
- **Interoperates**  
with other Intel parallel programming tools and libraries
- **Incremental**  
allows selective and targeted modification of existing code bases
- **Expressive**  
syntax oriented to application experts
- **Safe by default**  
deterministic semantics avoid race conditions and deadlock by construction
- **Easy to learn**  
serially consistent semantics and simple interface leverage existing skills
- **Widely applicable**  
Generalized data parallel model applicable to many types of computations

# Performance

- **Scalable to large problems**  
manages data to directly address memory bottlenecks
- **Unified thread and vector parallelization**  
single specification targets multiple mechanisms
- **Elimination of modularity overhead**  
automatically fuses multiple operations
- **Wide *and* deep**  
developers can choose level of abstraction  
can drill down to detail if needed

# Portability

- **High-level**  
avoids dependencies on particular hardware mechanisms or architectures
- **ISA extension independent**  
common binary can exploit different ISA extensions transparently
- **Allows choice of deployment hardware today**  
including scaling to many cores
- **Allows migration and forward-scaling**  
will support future hardware roadmap

ISA: Instruction Set Architecture

# Productivity via a High Level of Abstraction

"Specify what to do, not how to do it!"



**Mathematical structure**

**Data organization**

*Where's my data race?*

*What caused that deadlock?*

*Why do I get different answers  
every time I run this?*

*How many threads should I use?*

*How big is my cache?*

*How do I deal with different ISAs  
and vector widths?*

*Where's the guy who originally  
wrote this thing – I can't figure  
out what the code is supposed  
to be computing!*

**Mathematical structure**

**Data organization**






*Goal: increasing the efficiency of the  
expert application developer*



# Get the Best of Both Worlds

Attribute	Productivity (math/scripting languages)	Performance (threads with intrinsic)	Productivity & Performance (Array Building Blocks)
<b>Readability</b> Clear and understandable notation	Provides excellent support	Provides little or no support	Provides excellent support
<b>Determinism</b> Output is always the same for a given input	Provides excellent support	Provides little or no support	Provides excellent support
<b>Correctness</b> Major sources of error are avoided	Provides excellent support	Provides little or no support	Provides excellent support
<b>Performance</b> Best absolute performance	Provides little or no support	Provides excellent support	Provides something in-between
<b>Scalability</b> Ability to take advantage of increased number of cores	Provides little or no support	Provides something in-between	Provides excellent support
<b>Maintainability</b> Easy to maintain	Provides excellent support	Provides little or no support	Provides excellent support

 Provides excellent support
  Provides something in-between
  Provides little or no support

# Intel® ArBB vs. Intel® SSE intrinsics

```
template<typename T>
void C1D(dense<T> a)
{
    dense<T> w = abs(a);
    dense<T> k = 1.0f / (1.0f + 0.223646149f);
    w = 0.39894228f * k;
    0.355564737f * k * k;
    1.38447373f * k * k * k;
    1.621255978f * k * k * k * k;
    1.330274429f * k * k * k * k * k;

    w = w * sqrt_2 * exp(11 * -0.5f);
    w = select<T>(x > 1.0f, w);
    return w;
}

template<typename T>
void C2D(dense<T> a, dense<T> b)
{
    dense<T> c = b * a;
    dense<T> &res = c;
    dense<T> s = sqrt_2 * sqrt_2;
    dense<T> d = ln(a) + (r + v * 0.5f) / (s * sqrt_value);
    dense<T> e = sqrt_2 * sqrt_2;
    result = exp((d * r - c) / (d * C1D(s) + e)) / (1.0f - C1D(s));
}
```

- **vectorized**
- **threaded**
- **machine independent**

- **vectorized**
- **threaded**
- **machine independent**

# SS

```

type: strike, dtype: rate, dtype: volatility
dtype: time, int, dtype: float64))

int i
//local private working variables for the calculation
double cStrikePrice;
double hStrikeRate;
double aRiskAverseRate;
double aVolatility;
double aTime;
double aKerfTime;

__MMI_ALIGNED(dtype log(basis*NC0))
__MMI_KerfTerm;
__MMI_d0_d02;
__MMI_aPowerTerm;
__MMI_d0;

__MMI_ALIGNED(dtype d1(SMAD_WDTHTE;
__MMI_ALIGNED(dtype d2(SMAD_WDTHTE;
__MMI_ALIGNED(dtype FutureValue(SMAD_WDTHTE;
__MMI_ALIGNED(dtype NotKerf(SMAD_WDTHTE;
__MMI_ALIGNED(dtype NotKerf2(SMAD_WDTHTE;
__MMI_ALIGNED(dtype NegKerf4(SMAD_WDTHTE;
__MMI_ALIGNED(dtype NegKerf42(SMAD_WDTHTE;

cStrikePrice = __MMI_LOAD(sStrikePrice);
cStrikeRate = __MMI_LOAD(cStrikeRate);
aRiskAverseRate = __MMI_LOAD(aRiskAverseRate);
aVolatility = __MMI_LOAD(aVolatility);
aTime = __MMI_LOAD(aTime);

aKerfTime = __MMI_SORT(aTime);

for (i=0; i<SMAD_WDTH+1;i++)
{
log(basis[i]) = log(sStrikePrice[i]/strike[i]);

aKerfTerm = __MMI_LOAD(aKerf(basis[i]));

aPowerTerm = __MMI_ALIGNED(aVolatility*aVolatility);
__MMI_ALIGNED(aPowerTerm);
//avoid overflow in aPowerTerm, __MMI_SET(2.0,0.5);
//avoid overflow in aPowerTerm;
a1 = __MMI_SUB(aRiskAverseRate,aPowerTerm);
a2 = __MMI_LOAD(aKerfTerm);
a3 = __MMI_LOAD(aKerf2Term);
a4 = __MMI_LOAD(aKerfTime);
a5 = __MMI_LOAD(aVolatility);
a6 = __MMI_LOAD(aTime);
//if 1, 101 or 5, 101 or 5, an approximation is not too corruptive a02, but to derive it from a01
__MMI_LOAD(a02);
a02 = __MMI_LOAD(a02);
a02 = __MMI_LOAD(a02);
__MMI_STORE(a01);
__MMI_STORE(a02);
}

CNDIF(notKerf < 1;
CNDIF(notKerf2 < 2;

for (i=0; i<NC0; i++) {
FutureValue[i] = strike[i] * (exp(-a1*(a2+i*time[i])));
NegKerf4[i] = (-a3*a4+i*a5);
NegKerf42[i] = (-a3*a4+i*a5);
OptionPrice[i] = (FutureValue[i] + NegKerf42[i]) * (a3*a4+i*a5);
NegKerf42[i];
}

void setStrikeScholens(option_price,
int num_options,
dtype strikePrice,
dtype strikeRate,
dtype time,
dtype volatility,
dtype timeTime)
{
for (i=0; i<num_options; i++) {
// Calling main function to calculate option value based on Black &
Scholens'
// equation
BlackScholensCallorPut(d(option_price[i], NC0, a1(strikePrice[i]),
&strikeRate[i],
&aTime[i], &aVolatility[i], &aTimeTime[i], NAL7(a3(a4+i*a5)*(-i*a5))
)
}
}

```

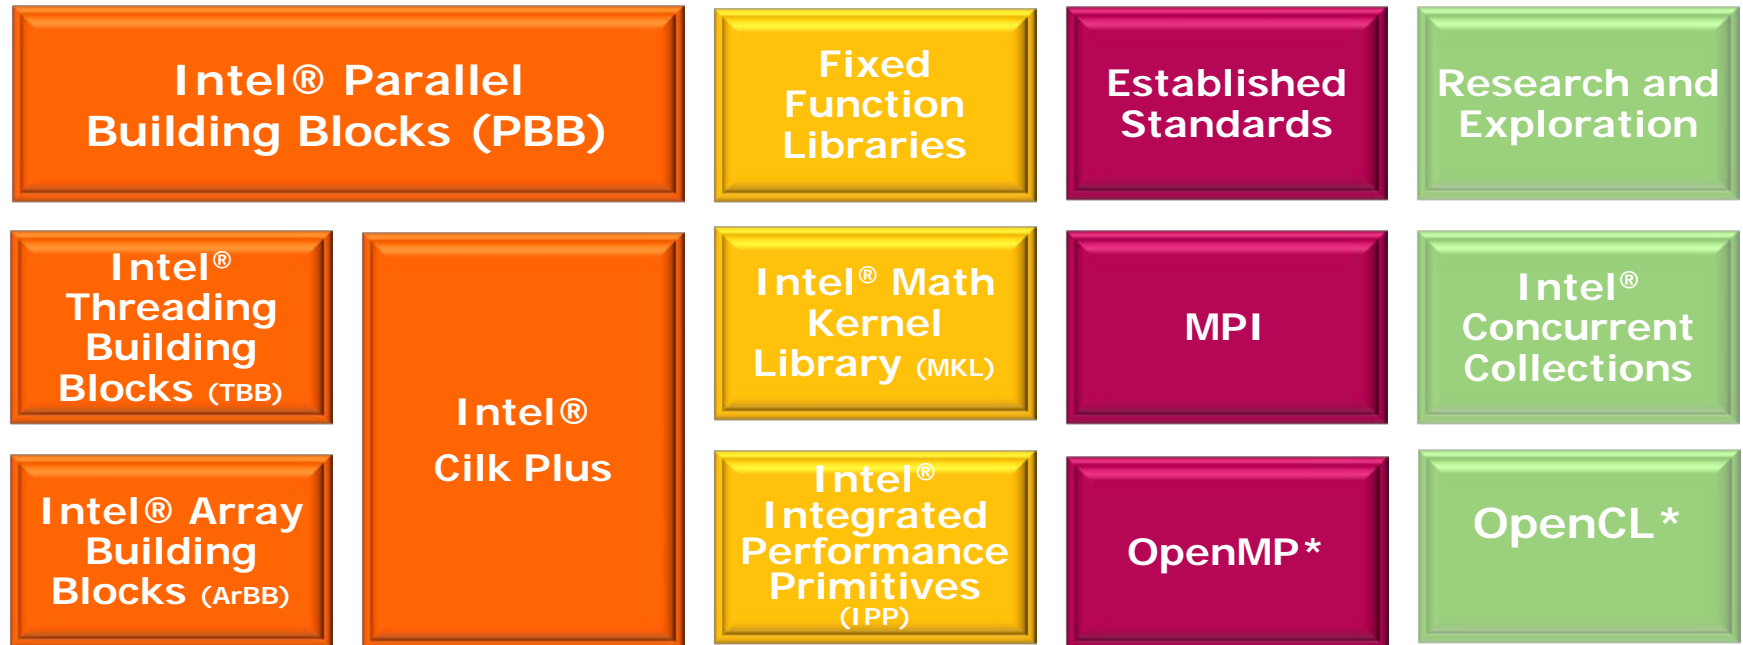
# 186 lines

- **vectorized**
- *not threaded*
- *machine dependent*



Copyright © 2010, Intel Corporation. All rights reserved.  
\*Other brands and names are the property of their respective owners.

# Intel's Family of Parallel Models



# Where do Customers Get them?

Launched Nov 9!

## Intel® Parallel Studio XE 2011



### Intel® Composer XE 2011

- Intel® C++ Compiler XE 12.0
- Intel® Fortran Compiler XE 12.0
- Intel® Parallel Debugger Extension
- Intel® Parallel Building Blocks (all)
- Intel® Math Kernel Library



### Intel® Inspector XE 2011



### Intel® Vtune™ Amplifier XE 2011

Windows: Integrates into Microsoft\* Visual Studio\* or stand-alone

Linux: Integrates into Eclipse CDT

1 Year Premier Support Renewable Annually

Launched Sep 2!

## Intel® Parallel Studio 2011



### Intel® Parallel Advisor 2011



### Intel® Parallel Composer 2011

- Intel® C++ Compiler XE 12.0
- Intel® Parallel Debugger Extension
- Intel® Parallel Building Blocks
- Intel® Threading Building Blocks
- Intel® Cilk™ Plus



### Intel® Parallel Inspector 2011



### Intel® Parallel Amplifier 2011

Windows: Integrates into Microsoft\* Visual Studio\*

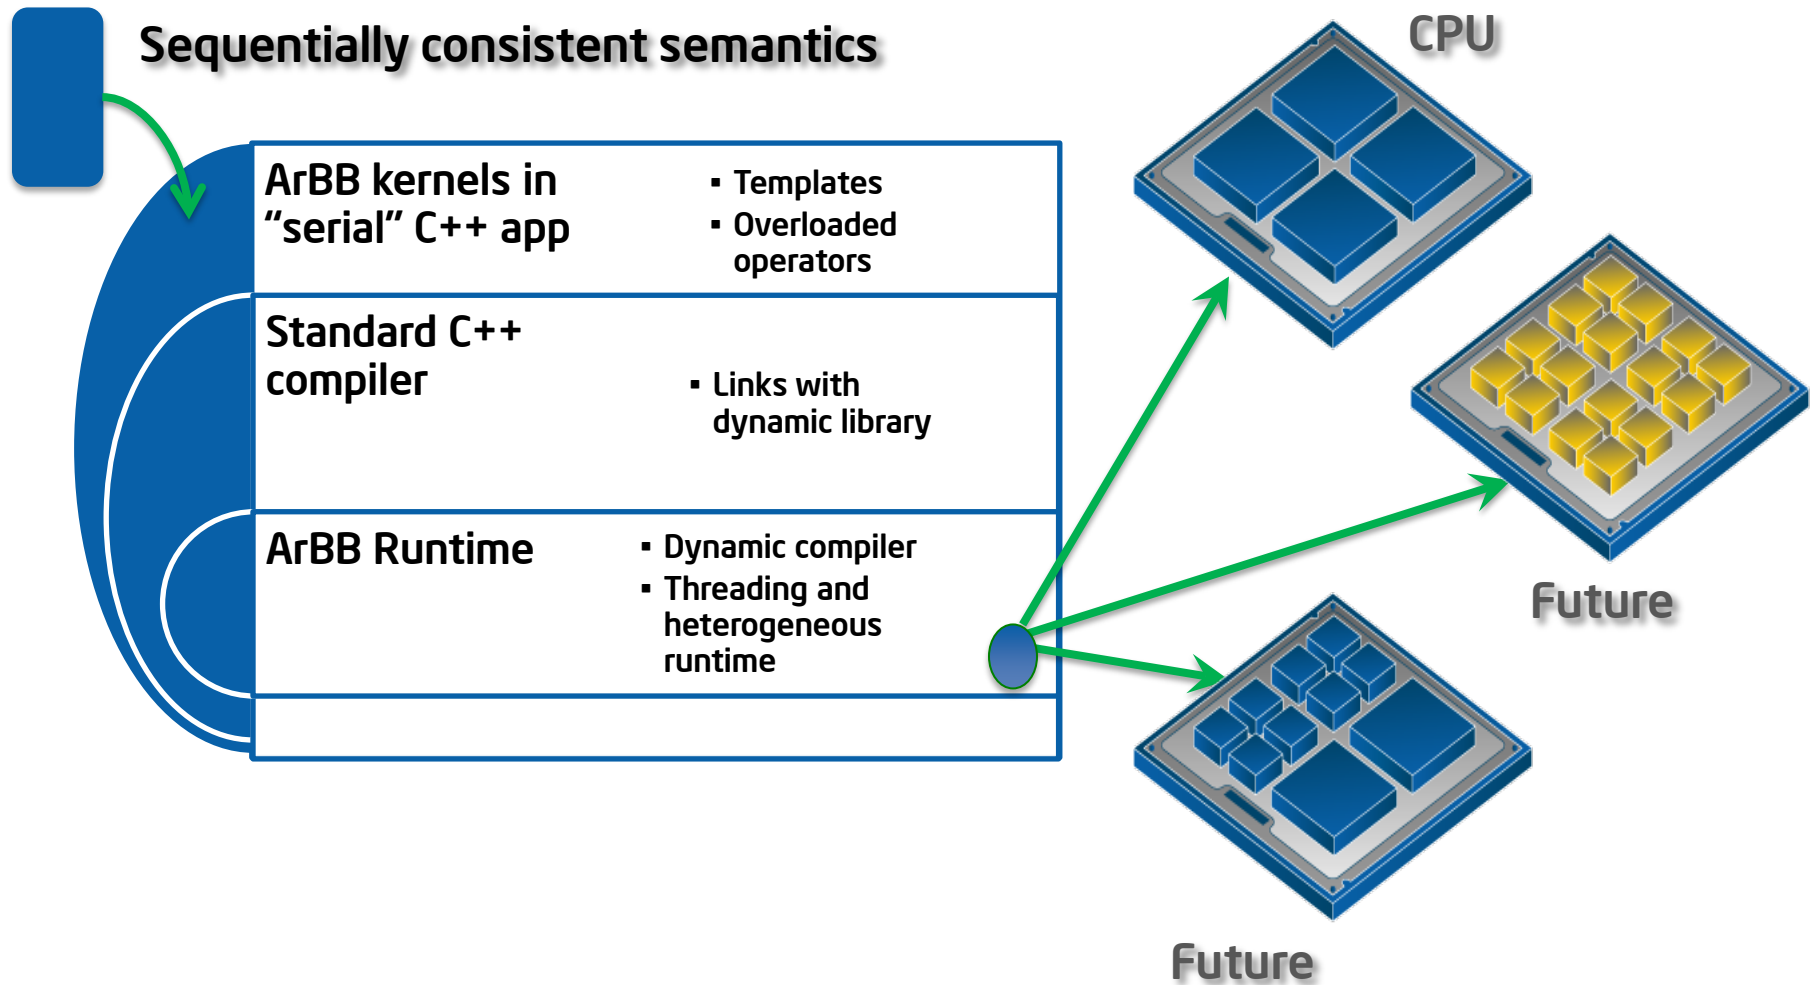
1 Year Premier Support Renewable Annually



# Levels of Parallelism

Grid	Group of clusters communicating through internet
Cluster	Group of computers communicating through fast interconnect
Node	Group of processors communicating through shared memory
Socket	Group of cores communicating through shared cache
Core	Group of functional units communicating through registers
Hyper-Threads	Group of thread contexts sharing functional units
Superscalar	Group of instructions sharing functional units
Pipeline	Sequence of instructions sharing functional units
Vector	Single instruction using multiple functional units

# How does it work?

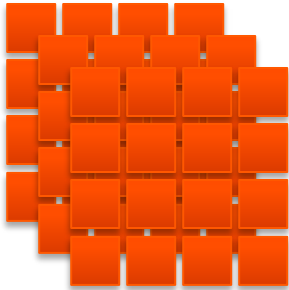


# Containers

## regular containers



*dense<T>*



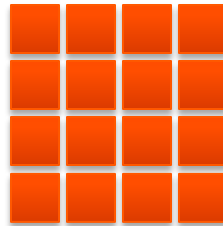
*dense<T,3>*



*array<...>*

*struct user\_type {..};*

*class user\_type { };*



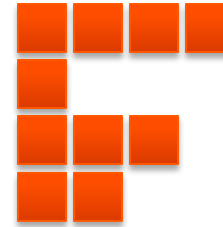
*dense<T, 2>*



*dense<array<...>>*

*dense<user\_type>*

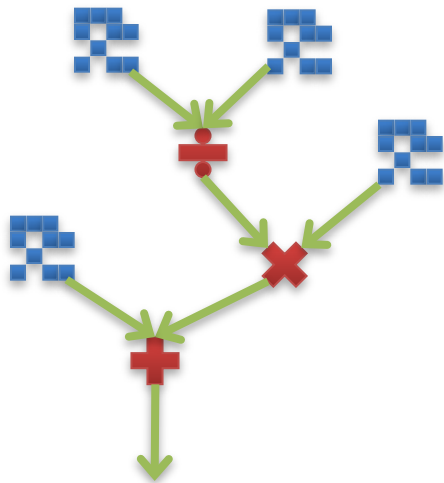
## irregular containers



*nested<T>*

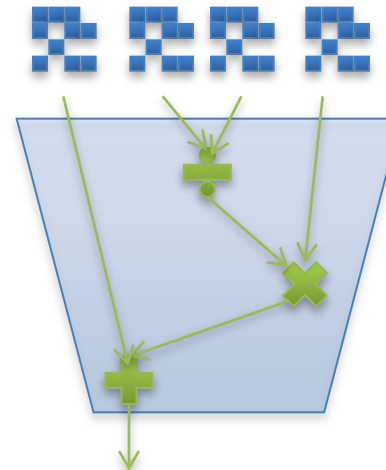
# Vector Processing *or* Scalar Processing

## Vector Processing



```
dense<f32> A, B, C, D;  
A = A + B/C * D;
```

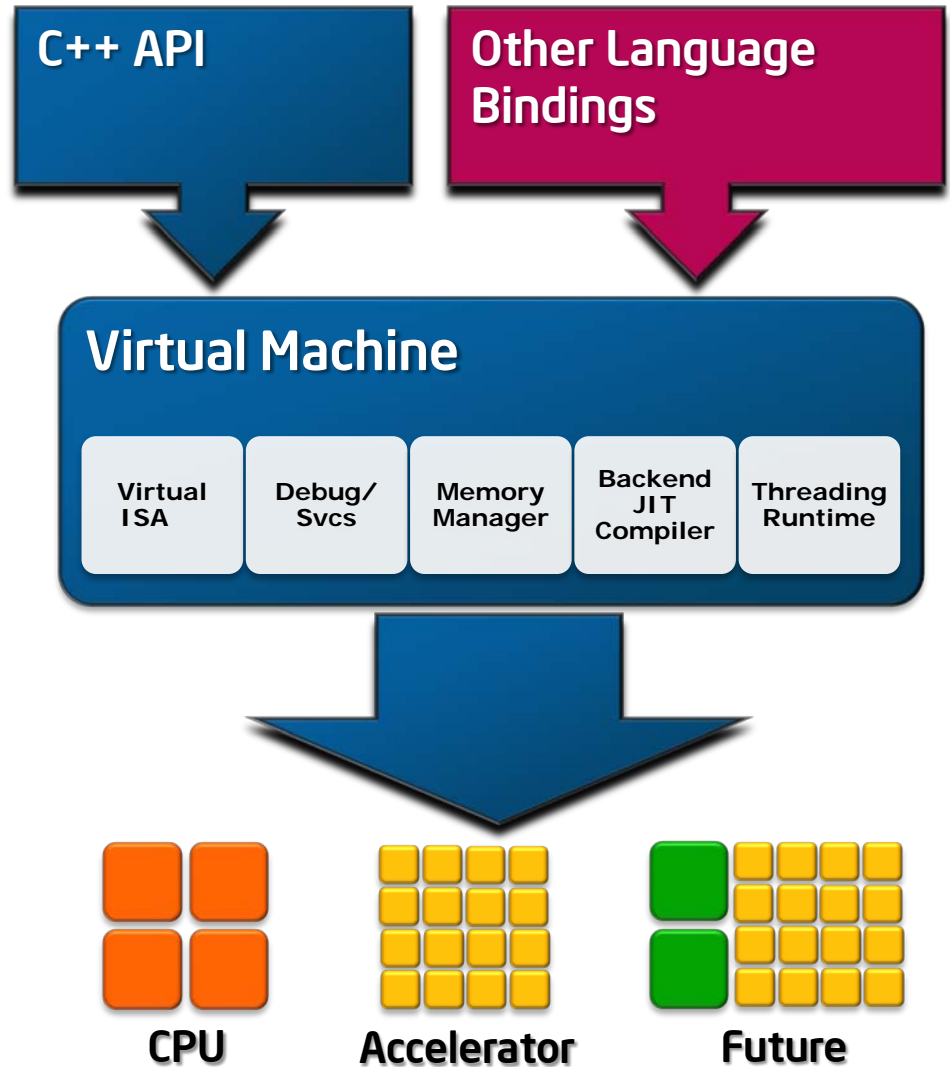
## Scalar Processing



```
void kernel(f32& a, f32 b, f32 c, f32 d) {  
    a = a + (b/c)*d;  
}  
...  
dense<f32> A, B, C, D;  
map(kernel)(A, B, C, D);
```

# Intel® ArBB Virtual Machine

- Generalized data-parallel programming model
- Supports wide variety of patterns and collections
- Supports explicit dynamic generation and management of code
- Implementation targets both threads and vector code
  - Machine independent optimization
  - Offload management
  - Machine specific code generation and optimizations
  - Scalable threading runtime



# Interface: The API as a Language

Syntax and semantics that extend C++

Adds parallel collection objects and methods to C++

- Uses standard C++ features (classes, simple templates, and operator overloading) to create new types and operators
- Sequences of API calls are fused and optimized by a JIT compiler

Works with standard C++ compilers

- Intel® C++ Compiler
- Microsoft\* Visual\* C++ Compiler
- GNU Compiler Collection\*

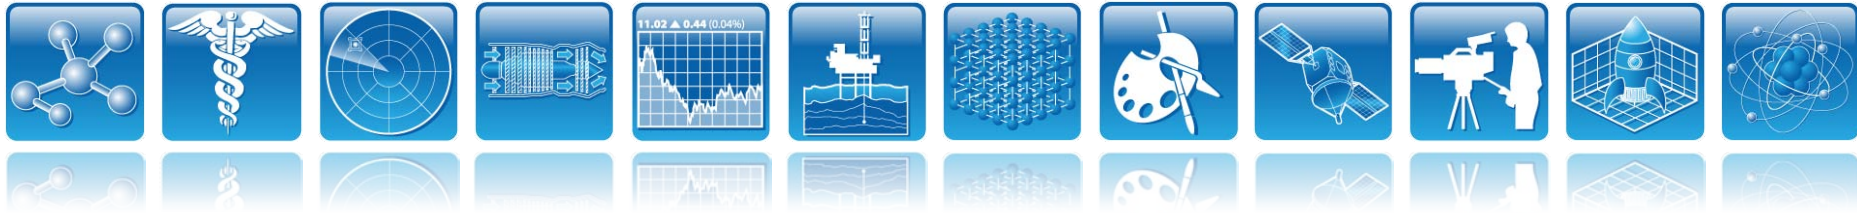
Express algorithms using mathematical notation

- Developers focus on *what to do*, not *how to do it*

Uses sequential semantics

- Developers do not use *threads*, *locks* or other lower-level constructs and can avoid the associated *complexity*
- *Programmers can reason and debug as if the program were serial.*

# What can it be used for?



## Bioinformatics

- Genomics and sequence analysis
- Molecular dynamics

## Engineering design

- Finite element and finite difference simulation
- Monte Carlo simulation

## Financial analytics

- Option and instrument pricing
- Risk analysis

## Oil and gas

- Seismic reconstruction
- Reservoir simulation

## Medical imaging

- Image and volume reconstruction
- Analysis and computer aided detection (CAD)

## Visual computing

- Digital content creation (DCC)
- Physics engines and advanced rendering
- Visualization
- Compression/decompression

## Signal and image processing

- Computer vision
- Radar and sonar processing
- Microscopy and satellite image processing

## Science and research

- Machine learning and artificial intelligence
- Climate and weather simulation
- Planetary exploration and astrophysics

## Enterprise

- Database search
- Business information

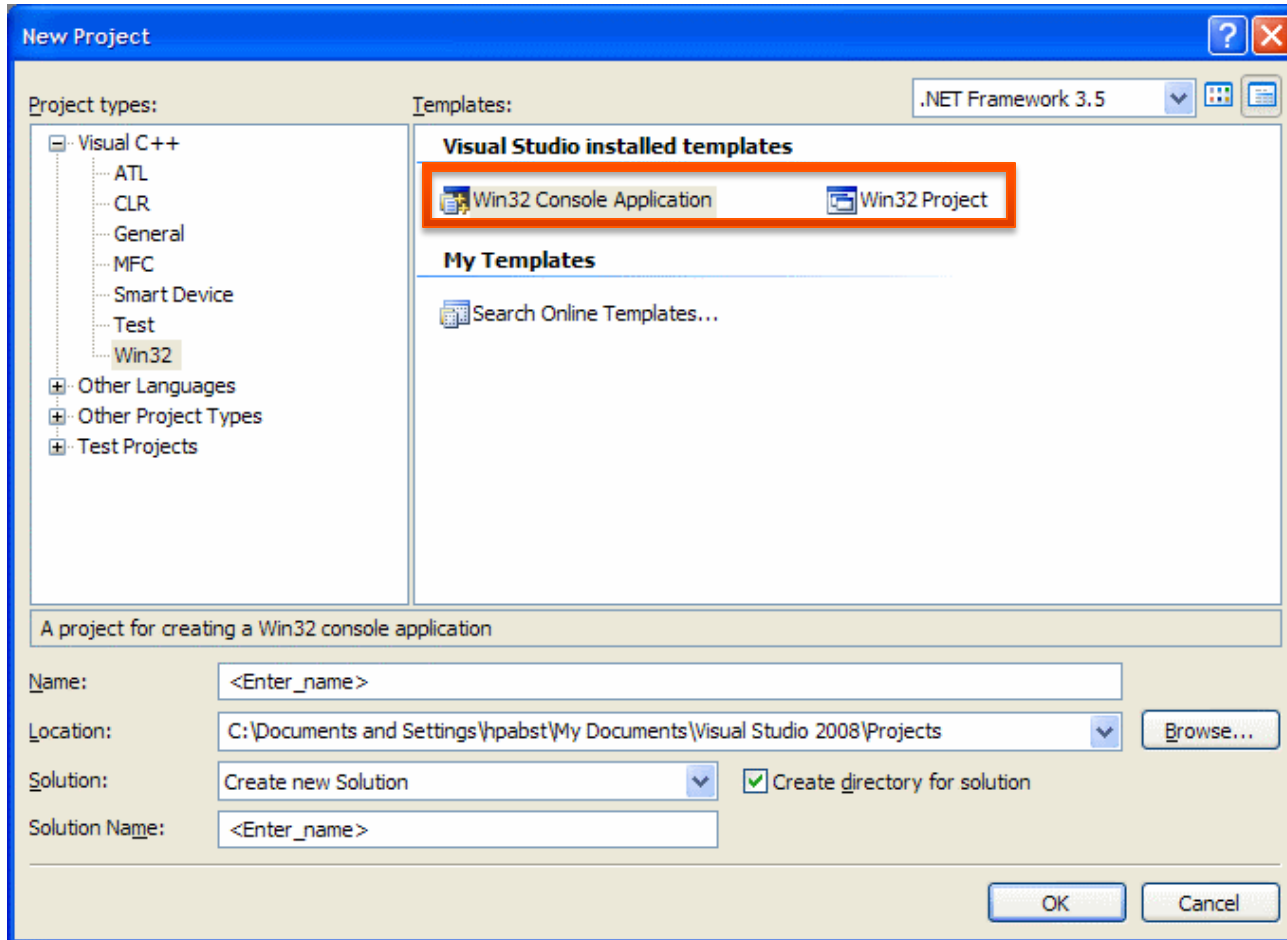


# Introduction to Intel® Array Building Blocks

How to add it to your project...

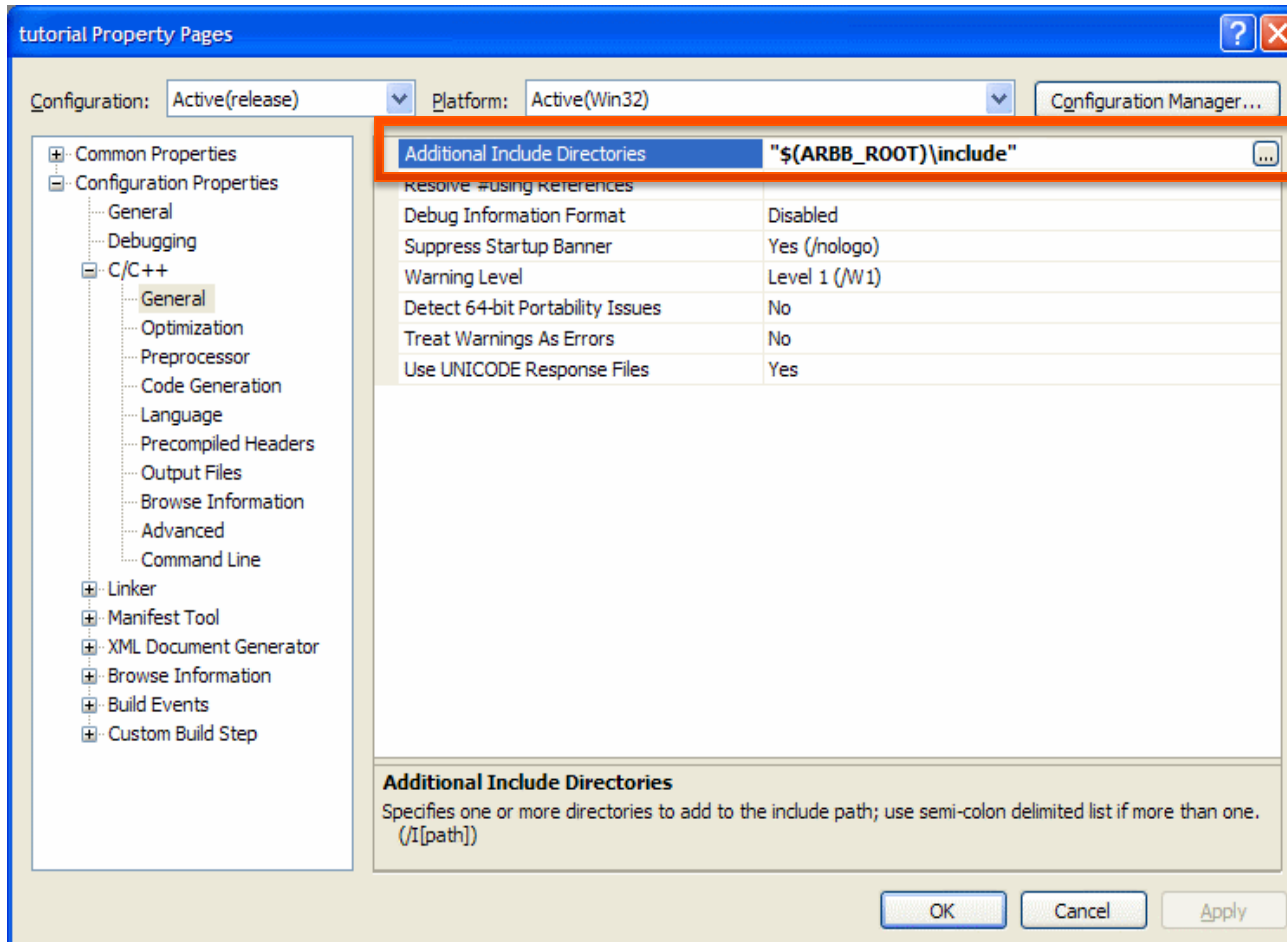


# Intel® ArBB in a Visual Studio\* Project



Screenshots taken from Microsoft\* Visual Studio 2008\*

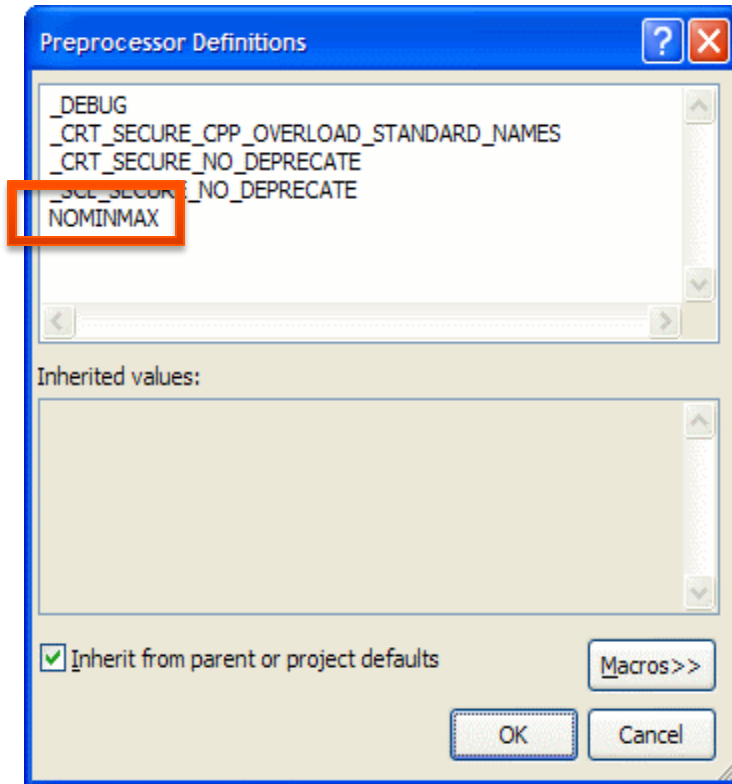
# Intel® ArBB in a Visual Studio\* Project



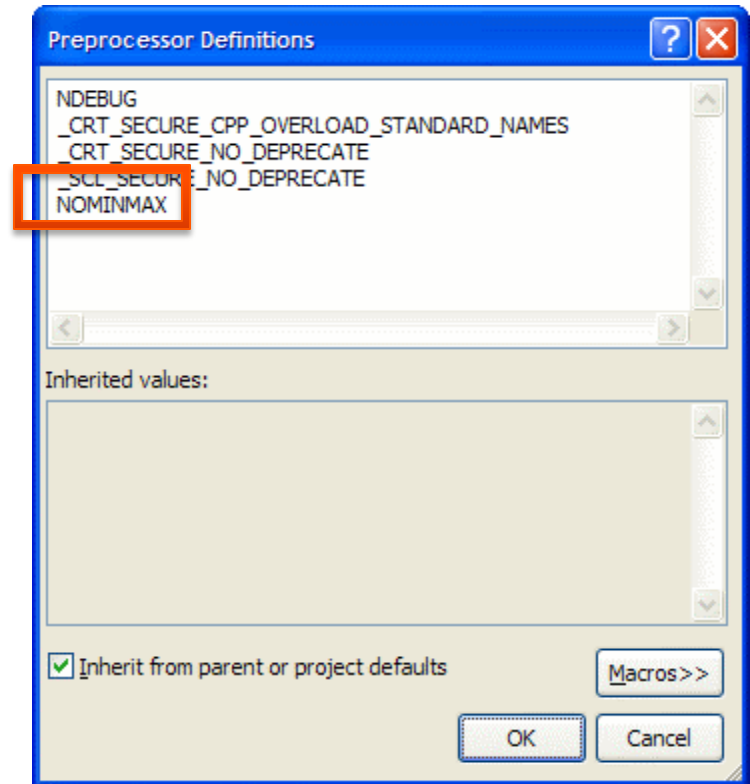
Screenshots taken from Microsoft\* Visual Studio 2008\*

# Including ArBB in a Visual Studio\* Project

## Debug Mode

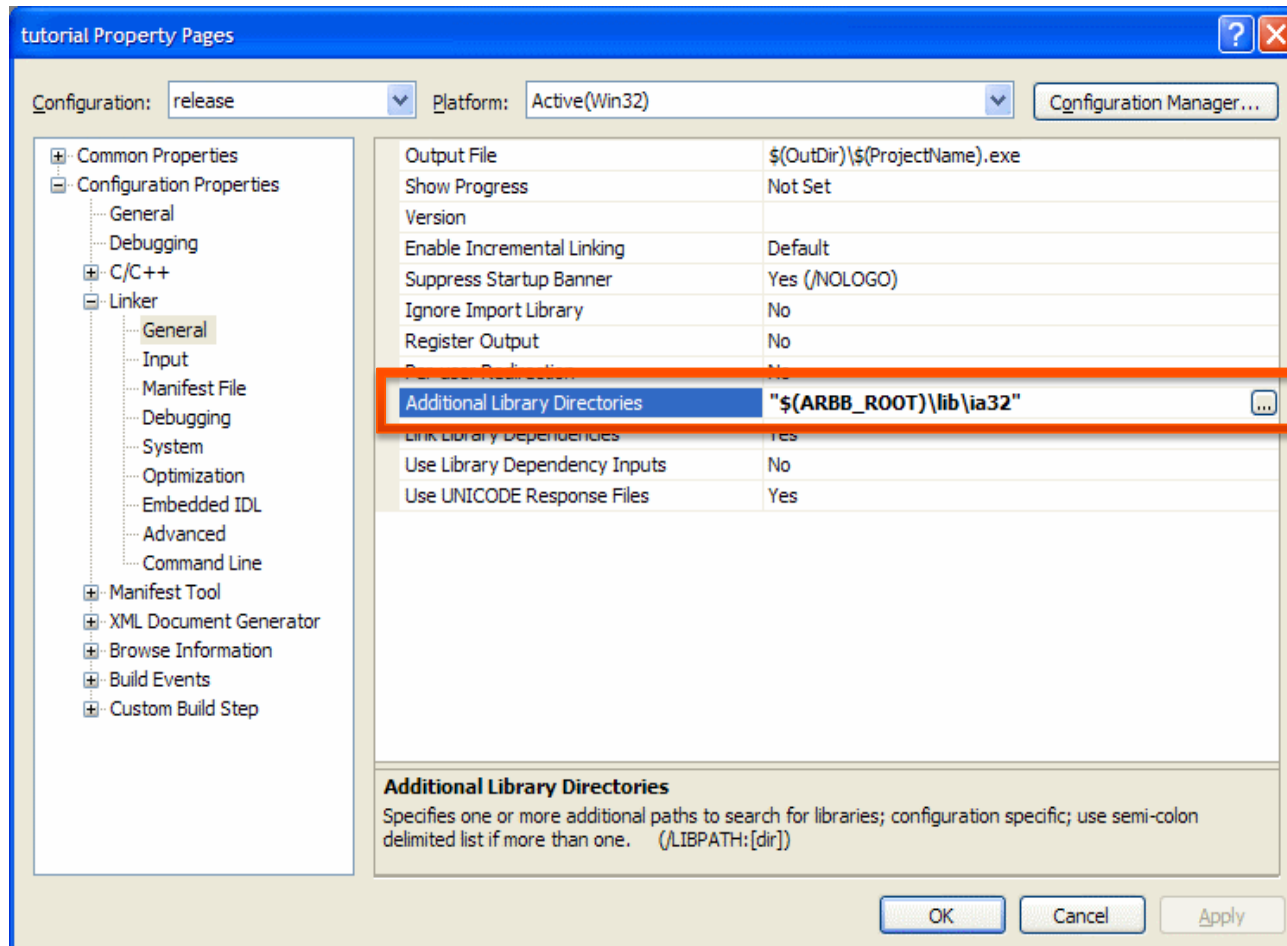


## Release Mode



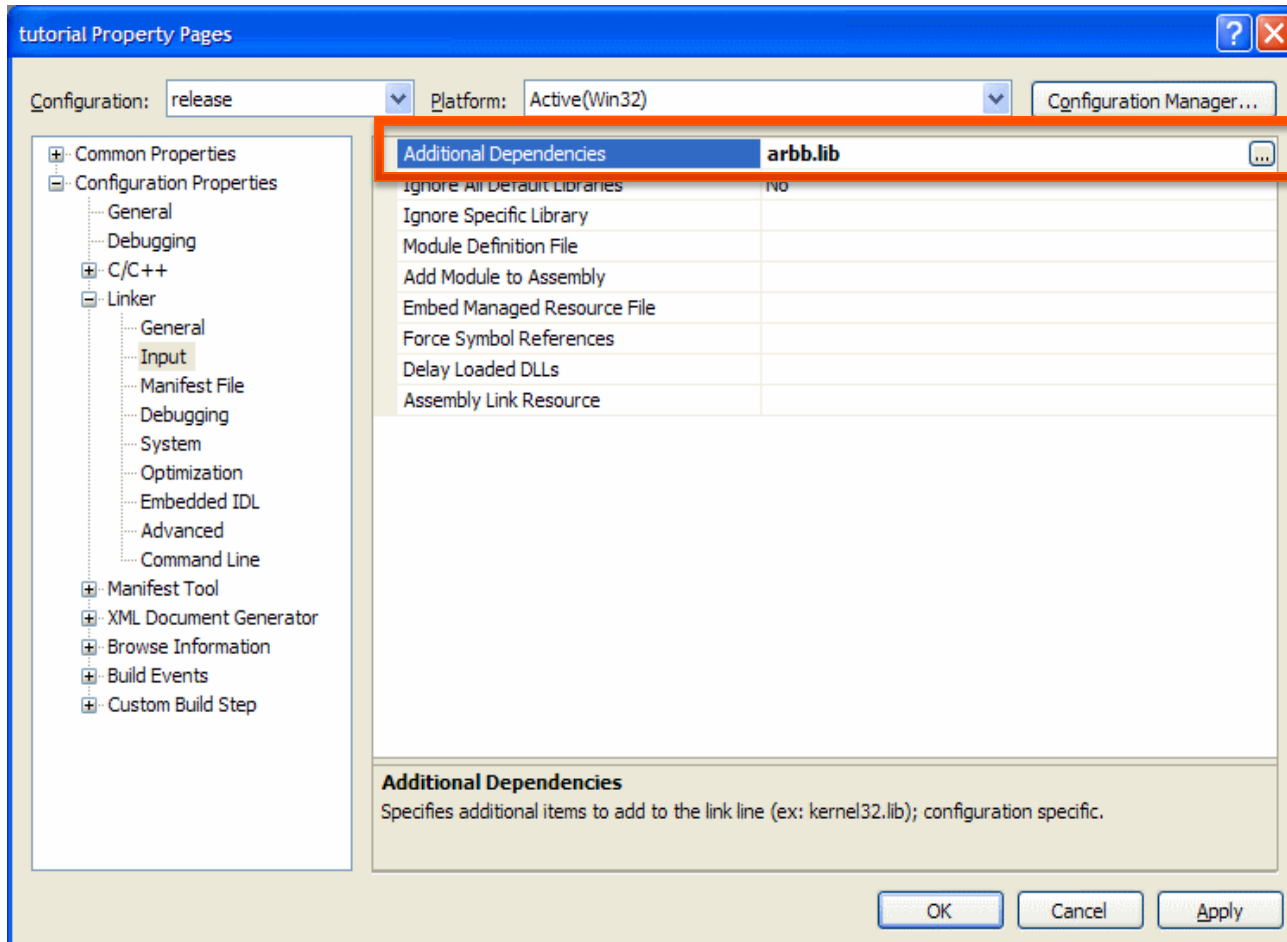
Screenshots taken from Microsoft\* Visual Studio 2008\*

# Intel® ArBB in a Visual Studio\* Project



Screenshots taken from Microsoft\* Visual Studio 2008\*

# Intel® ArBB in a Visual Studio\* Project



Screenshots taken from Microsoft\* Visual Studio 2008\*

# Intel® ArBB in a Makefile-based Project

- **Make available ArBB include (header) files:**
  - `-I/opt/intel/arbb/include`  
(modify compiler search path for include files)
- **Make available ArBB libraries**
  - `-L/opt/intel/arbb/lib/{ia32,intel64}`  
(modify linker search path for libraries)
- **Include ArBB libraries in linker process**
  - `-larbb -ltbb`

# Using the Intel® ArBB API

- Include the definitions

```
#include <arbb.hpp>
```

- Import the namespace or specific identifiers

```
using namespace arbb;
```

```
using namespace arbb::add_reduce;
```

- Good practice:

- To not pollute the name spaces, restrict scope of “using” statement as much as possible, especially in headers
- Selectively include ArBB names only if used

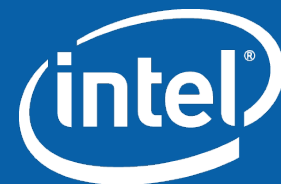
# Code Skeleton for Intel® ArBB Applications

- Use the following code skeleton for ArBB applications

```
int main(int argc, char* argv[]) {  
    int ret_code;  
    try {  
        // call into ArBB code  
        ret_code = EXIT_SUCCESS;  
    }  
    catch(const std::exception& e) {  
        ret_code = EXIT_FAILURE;  
    }  
    catch(...) {  
        cerr << "Error: Unknown exception caught!" << endl;  
        ret_code = EXIT_FAILURE;  
    }  
    return ret_code;  
}
```

- ArBB indicates runtime errors through standard C++ exceptions
- Existing top-level entry points do not need to change if they already catch std::exception





# Introduction to Intel® Array Building Blocks

Programming Constructs and Data Types

# Overall Syntax Conventions

- **All Identifiers are lower-case with underscores**
  - `some_type`
  - `some_class::some_member_function()`
- **Chosen to align with C++ standard library conventions**

# Intel® ArBB Constructs

- **Scalar types**
  - Equivalent to primitive C++ types
- **Vector types**
  - Parallel collections of (scalar) data
- **Operators**
  - Scalar operators
  - Vector operators
- **Functions**
  - User-defined code fragments
- **Control flow constructs**
  - Conditionals, iteration, etc.
  - These are for *serial* control flow *only*
  - Vector operations and “map” are used for expressing parallelism

# Scalar types

- **Scalar types provide equivalent functionality to the scalar types built into C/C++**

Types	Description	C++ equivalents
f32, f64	32/64 bit floating point number	float, double
i8, i16, i32, i64	8/16/32 bit signed integers	char, short, int
u8, u16, u32, u64	8/16/32 bit unsigned integers	unsigned char/short/int
boolean	Boolean value (true or false)	bool
usize, isize	Signed/unsigned integers sufficiently large to store addresses	size_t (eqv. usize)

# Scalar Types

## Use scalar types for ArBB scalar computation

```
i32 int_scalar; // a scalar 32-bit integer value  
f32 fp_scalar = (f32)int_scalar; // cast a scalar to new type
```

## Casting to/from C/C++ types

```
float f = (float)fp_scalar; // NOT supported  
f32 fp_scalar2(f); // immediate copy  
f32 fp_scalar3 = f; // immediate copy  
float x = value(fp_scalar); // retrieve value
```

## Constant values are supported (types must match)

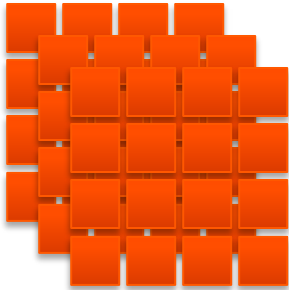
```
f32 fp_scalar = (f32)int_scalar + 0.5f;  
f32 r = 2.0f;  
fp_scalar = 3.14f * r * r;
```

# Containers

## regular containers



*dense<T>*



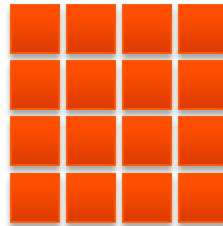
*dense<T,3>*



*array<...>*

*struct user\_type {..};*

*class user\_type { };*



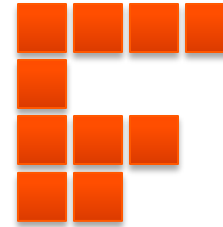
*dense<T, 2>*



*dense<array<...>>*

*dense<user\_type>*

## irregular containers



*nested<T>*

# Dense Containers

```
template<typename T, std::size_t D = 1>
class dense;
```

- **This is the equivalent to `std::vector` or C arrays**
- **Dimensionality is optional, defaults to 1**

Property	Restrictions	Can be set at
Element type	Must be an ArBB scalar or user-defined type	Compile time
Dimensionality	1, 2, or 3	Compile time
Size	Only restricted by free memory	Runtime

# Declaration and Construction

Declaration	Element type	Dimensionality	Size
dense<f32> a1;	f32	1	0
dense<f32, 1> a2;	f32	1	0
dense<i32, 2> b;	i32	2	0, 0
dense<f32> c(1000);	f32	1	1000
dense<f32> d(c);	f32	1	1000
dense<i8, 3> e(5, 3, 2);	i8	3	5, 3, 2



# Operations on dense Containers

- **All scalar operations can be applied element-wise**
  - Arithmetic and bit operations, transcendentals, etc.
- **Additionally provides container operations:**
  - Indexing, e.g. `operator[]`
  - Reordering, e.g. `shift()`, `section()`
  - Reductions, e.g. `sum()`, `any()`, `all()`
  - Prefix sums, packs, and other data-parallel primitives
  - Property access, e.g. `num_rows()`
- **Most of these operations run in parallel**
  - For example, if you add two dense containers together, all the individual additions can run in parallel

# Moving Data into and out of Containers

- Dense containers provide two ways to access data:
  - Iterators
    - `read_only_range` iterator to read from the container
    - `write_only_range` iterator to write into the container
    - `read_write_range` iterator to write/read a container
  - Binding
    - On construction, dense containers can be *bound* (associated) to a particular data location
    - Moves data into and out of that location when required

# Creating “dense” Containers

## Declaration of a dense container:

*// create an empty container whose values will be assigned later*

```
dense<f32> temp;
```

## vector objects of different base types cast into each other:

```
dense<i32> vi = ...;
```

```
dense<f32> v = (dense<f32>)vi;
```

# Filling “dense” Containers

*// request write-only access to container*

```
dense<f32> a(1024);  
range<f32> range_a = a.write_only_range();  
std::fill(range_a.begin(),  
           range_a.end(),  
           static_cast<f32>(1));
```

*// request read/write access to container*

```
dense<f32> b(1024);  
range<f32> range_b = b.read_write_range();  
std::fill(range_b.begin(),  
           range_b.end(),  
           static_cast<f32>(2));
```

# Fixed-size Arrays

- Typical usages: pairs of data, RGBA data, CYMK data, etc.
- Use `std::array` look-a-like
  - `std::array` is a C++ TR1/C++0x type
  - Will support `std::array` operations
  - You can manipulate with element-wise, horizontal, swizzling, and other utility operations

---

```
array<f32, 3> p1, p2, p3;
```

```
f32 r = p1[0];
```

```
p1 = p2 + p3;
```

```
f32 sum_p1 = sum(p1);
```

```
p1 = cat(p2, p3);
```

```
// std::array operations
```

```
// element-wise operations
```

```
// horizontal operations
```

```
// utility operations
```

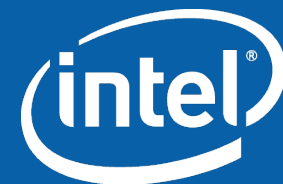
# Structured Types

- **C++ classes and structures can be used relatively normally within ArBB**
  - Requires that primitive types be classes in ArBB types (f32, etc.)
  - Supports member functions, class members, overloaded operators, etc.
  - However, virtual functions and pointers are resolved during “capture time” only
  - Overloaded operators are automatically lifted over collections
  - Lifting member functions over collections requires an additional declaration (a macro is provided to help with this)

# Structure/Class Example

```
class my_class {
public:
    my_class(f32 location, i32 count);
    my_class operator+(const my_class& other) {
        return my_class(location + other.location,
                          max(count, other.count));
    }
    // other code...
private:
    f32 m_location;
    i32 m_count;
};

dense<my_class> A, B, C;
A = B + C; // This will use the user-defined operator+!
my_class m = A[5]; // Other interactions work naturally.
```



**BREAK**



